

Customer Number: 23624

[1] MOVABLE TAP FINITE IMPULSE RESPONSE FILTER
[2] CROSS-REFERENCE TO RELATED APPLICATION

[3] This application is a continuation-in-part application of U.S. Serial No. 09/678,728, filed on October 4, 2000 and entitled "Movable Tap Finite Impulse Response Filter", the contents of which are incorporated herein by reference.

[4] BACKGROUND OF THE INVENTION

[5] Field Of The Invention

[6] The present invention relates to a finite impulse response filter, and particularly to such a filter in which a delay in a portion thereof has an adjustable or selectable delay period, and to an echo canceller and an Ethernet transceiver including such an FIR filter.

[7] Description Of The Related Art

[8] Finite impulse response (FIR) filters are extremely versatile digital signal processors that are used to shape and otherwise to filter an input signal so as to obtain an output signal with desired characteristics. FIR filters may be used in such diverse fields as Ethernet transceivers, read circuits for disk drives, ghost cancellation in broadcast and cable TV transmission, channel equalization for communication in magnetic recording, echo cancellation, estimation/prediction for speech processing, adaptive noise cancellation, etc. For example, see U.S. Patent Nos. 5,535,150; 5,777,910; and 6,035,320, the contents of each of which are incorporated herein by reference. Reference is also made to the following publications: "An adaptive Multiple Echo Canceller for Slowly Time Varying Echo Paths," by Yip and Etter, IEEE Transactions on Communications, October 1990; "Digital Signal Processing", Alan V. Oppenheim, et al., pp. 155-163; AA 100MHz Output Rate Analog-to-Digital Interface for PRML Magnetic-Disk Read Channels in 1.2um CMOS®, Gregory T. Uehara and Paul R. Gray, ISSCC94/Session 17/Disk-Drive Electronics/ Paper FA 17.3, 1994 IEEE International Solid-State Circuits Conference, pp. 280-281; "72Mb/S PRML Disk-Drive Channel Chip with an Analog Sampled Data Signal Processor", Richard G. Yamasaki, et al., ISSCC94/Session 17/Disk-Drive Electronics/Paper FA 17.2, 1994 IEEE International

Customer Number: 23624

Solid-State Circuits Conference, pp. 278, 279; "A Discrete-Time Analog Signal Processor for Disk Read Channels", Ramon Gomez, et al., ISSCC 93/Session 13/Hard Disk and Tape Drives/Paper FA 13.1, 1993

[9] ISSCC Slide Supplement, pp. 162, 163, 279, 280; and AA 50MHz 70 mW 8-Tap Adaptive Equalizer/Viterbi Sequence Detector in 1.2 μ m CMOS®, Gregory T. Uehara, et al. 1994 IEEE Custom Integrated Circuits Conference, pp. 51-54, the contents of each being incorporated herein by reference.

[10] Typically, an FIR filter is constructed in multiple stages, with each stage including an input, a multiplier for multiplication of the input signal by a coefficient, and a summer for summing the multiplication result with the output from an adjacent stage. The coefficients are selected by the designer so as to achieve the filtering and output characteristics desired in the output signal. These coefficients (or filter tap weights) are often varied, and can be determined from a least mean square (LMS) algorithm based on gradient optimization. The input signal is a discrete time sequence which may be analog or digital, while the output is also a discrete time sequence which is the convolution of the input sequence and the filter impulse response, as determined by the coefficients.

[11] With such a construction, it can be shown mathematically and experimentally that virtually any linear system response can be modeled as an FIR response, as long as sufficient stages are provided. Because of this feature, and the high stability of FIR filters, such filters have found widespread popularity and are used extensively.

[12] One problem inherent in FIR filters is that each stage requires a finite area on an integrated circuit chip. Additional area is required for access to an external pin so as to supply the multiplication or weighting coefficient for that stage. In some environments, the number of stages needed to provide desired output characteristics is large. For example, in Gigabit Ethernet applications it is preferred that every 8 meters of cable length be provided with 11 stages of FIR filter. In order to cover cable lengths as long as 160 meters, 220 FIR stages should be provided. In such environments, chip area on the integrated circuit is largely monopolized by the FIR stages.

[13] Moreover, each FIR stage requires a finite amount of

Customer Number: 23624

power and generates a corresponding amount of heat. Particularly where a large number of stages is needed, such power requirements become excessive and require significant mechanical adaptations to dissipate the heat.

[14] The inventors herein have recently recognized that in some environments, not all stages of an FIR contribute significantly to the output. Figure 1, for example, is a waveform showing signal amplitude versus time in an Ethernet echo cancellation application, where time (on the horizontal axis) is expressed in delay units for an FIR filter. The waveform shown in Figure 1 represents an Ethernet transmission and its echo (or, reflection). As seen in Figure 1, the waveform includes the near end echo at region 1, followed by a relatively quiet period in region 2, a relatively negligible signal at region 3, and the far end echo at region 4. One use of an FIR filter in such an Ethernet environment is to cancel the echo so as to distinguish more clearly between incoming signals and simple reflections of transmitted signals. However, the relatively negligible signal at region 3 contributes very little to the overall output of the FIR filter. The reason for this is that, whatever value of coefficients are set at the stages corresponding to region 3, those coefficients will be multiplied by a value which is approximately zero. Thus, contributions of those signals to the FIR output will be negligible, especially compared to regions 1, 2 or 4.

[15] The inventors have considered simplifying the selection of coefficients by setting the coefficients corresponding to region 3 to zero, which would result in simpler algorithms needed to select coefficients. However, even with zeroed coefficients, the stages corresponding to region 3 still exist on the integrated circuit chip, stealing valuable surface area and power, and generating unwanted heat.

[16] SUMMARY OF THE INVENTION

[17] According to a first aspect of the present invention, an FIR filter is provided comprising a coefficient generator to generate first and second coefficients, a first control conductor, and a second control conductor. A controller is coupled to a first end of the first control conductor and a first end of the second control conductor. A shared wiring is provided having its first end being coupled to the coefficient generator. A first memory is coupled to a second end of the shared wiring and

Customer Number: 23624

coupled to a second end of the first control conductor to store the first coefficient in response to the controller.

A first multiplier is responsive to the first coefficient stored in the first memory and the input, and a first delay circuit is responsive to an input. A second memory is coupled to the second end of the shared wiring and coupled to a second end of the second control conductor to store the second coefficient in response to the controller, and a second multiplier is responsive to the second coefficient stored in the second memory and the first delay element.

[18] According to a second aspect of the present invention, an FIR filter apparatus having N taps, N being a positive integer of at least two, is provided comprising a coefficient generator to generate N coefficients, one for each of the N taps. A shared wiring is provided responsive to an output of the coefficient generator. N memories is provided, each being responsive to the shared wiring to store a respective one of the N coefficients. An FIR filter comprises N filter stages, each stage being responsive to one of the N coefficients stored in a corresponding one of the N memories.

[19] According to a third aspect of the present invention, an FIR filter apparatus comprises a coefficient generator to generate first and second coefficients. A shared wiring is provided having its first end coupled to the coefficient generator. A first memory is coupled to a second end of the shared wiring to store the first coefficient in response to a selector. A first multiplier is responsive to the first coefficient stored in the first memory and an input. A first delay circuit responsive to the input, and a second memory is coupled to the second end of the shared wiring to store the second coefficient in response to the selector. A second multiplier is provided responsive to the second coefficient stored in the second memory and the first delay element.

[20] According to a fourth aspect of the present invention, an FIR filter apparatus having N taps, N being a positive integer of at least two, comprises a coefficient generator to generate N coefficients, one for each of the N taps. A shared wiring is responsive to an output of the coefficient generator. N memories are provided, each of the memories being responsive to the shared wiring to store a respective one of the N coefficients in response to a selector. An FIR filter comprises N filter stages, each one of the N filter stages being responsive to one of the N coefficients stored in a corresponding one of the N memories.

[21] According to a fifth aspect of the present invention, an FIR filter apparatus comprises a coefficient generator means for generating first and second coefficients. A

Customer Number: 23624

controller means synchronizes the coefficient generator, and a first control conductor means transfers a first control signal from the controller means. A second control conductor means transfers a second control signal from the controller means, and a shared wiring means transfers the first and second coefficients from the coefficient generator means. An input means for inputting a signal, and a first memory means stores the first coefficient transferred by the shared wiring means in response to the first control signal transferred by the first control conductor means. A first multiplier means multiplies the first coefficient stored in the first memory means by the signal input to the input means, and a first delay means delays the signal input to the input means. A second memory means stores the second coefficient transferred by the shared wiring means in response to the second control signal transferred by the second control conductor means, and a second multiplier means multiplies the second coefficient stored in the second memory means by the signal delayed by the first delay means.

[22] According to a sixth aspect of the present invention, an FIR filter apparatus having N taps, N being a positive integer of at least two, comprises a coefficient generator means for generating N coefficients, one for each of the N taps. A shared wiring means is provided for transferring the N coefficients from the coefficient generator means, and N memory means are provided, each of the memory means being responsive to the shared wiring means for storing a respective one of the N coefficients. An FIR filter means for filters an input signal comprising N filter stages, each one of the N filter stages being responsive to one of the N coefficients stored in a corresponding one of the N memory means.

[23] According to a seventh aspect of the present invention, an FIR filter apparatus comprises a coefficient generator means for generating first and second coefficients. A shared wiring means is provided for transferring the first and second coefficients from the coefficient generator means. A first memory means storing the first coefficient transferred by the shared wiring means in response to a selector signal from a selector means. A first multiplier means multiplies the first coefficient stored in the first memory means by a signal input to an input means, and a first delay means delays the signal input to the input means. A second memory means stores the second coefficient transferred by the shared wiring means in response to the selector signal from the selector means, and a second multiplier means multiplies the second coefficient stored in the second memory means by the signal delayed by the first delay

Customer Number: 23624

means.

[24] According to an eighth aspect of the present invention, an FIR filter apparatus having N taps, N being a positive integer of at least two, comprises a coefficient generator means for generating N coefficients, one for each of the N taps. A shared wiring means is provided for transferring the N coefficients from the coefficient generator means, and a selector means generates a selection signal. N memory means are provided, each of the memory means stores a corresponding one of the N coefficients transferred by the shared wiring means in response to the selection signal from the selector means. An FIR filter means filters a signal and comprises N filter stages, each one of the N filter stages being responsive to one of the N coefficients stored in a corresponding one of the N memory means.

[25] According to a ninth aspect of the present invention, a method of filtering a signal comprises (a) generating first and second coefficients; (b) synchronizing the generation of the first and second coefficients from step (a); (c) transferring a first control signal from step (b); (d) transferring a second control signal from step (b); (e) providing a shared wiring for transferring the first and second coefficients; (f) inputting a signal; (g) storing the first coefficient transferred in step (e) in response to the first control signal transferred in step (c); (h) multiplying the first coefficient stored in step (g) by the signal input in step (f); (i) delaying the signal input in step (f); (j) storing the second coefficient transferred in step (e) in response to the second control signal transferred in step (d); and (k) multiplying the second coefficient stored in step (j) by the signal delayed in step (i).

[26] According to a tenth aspect of the present invention, a method of filtering a signal comprises (a) generating N coefficients; (b) providing a shared wiring for transferring the N coefficients generated in step (a); (c) storing the N coefficients transferred in step (b); (d) filtering an input signal responsive to the N coefficients stored step (c); and synchronizing step (a) and step (c).

[27] According to an eleventh aspect of the present invention a method of filtering a signal comprises (a) generating first and second coefficients; (b) providing shared wiring for transferring the first and second coefficients generating in step (a); (c) inputting a signal;

[28] (d) providing a selector signal; (e) storing the first coefficient transferred by step (b) in response to the selector signal from step (d); (f) multiplying the first coefficient stored in step (e) by the signal in step (c); (g) delaying the signal input in step (c); (h)

Customer Number: 23624

storing the second coefficient transferred by step (b) in response to the selector signal from step (d); and multiplying the second coefficient stored in step (h) by the signal delayed in step (g).

[29] According to a twelfth aspect of the present invention a method of filtering a signal comprises (a) generating N coefficients; (b) providing a shared wiring for transferring the N coefficients from step (a); (c) generating a selection signal; (d) storing the N coefficients transferred in step (b) in response to the selection generated in step (c); (e) filtering a signal responsive to the N coefficients stored in step (d); and (f) synchronizing step (a) with step (d).

[30] This brief summary has been provided so that the nature of the invention may be understood quickly. A more complete understanding of the invention can be obtained by reference to the following detailed description of the preferred embodiments in connection with the attached drawings.

[31] BRIEF DESCRIPTION OF THE DRAWINGS

[32] Figure 1 is a view showing a channel response waveform over copper cable in an Ethernet environment, including near end echo and far end echo due to reflection.

[33] Figure 2 is a functional block diagram showing an Ethernet transceiver including a transmit side and a receive side, and in which an echo canceller thereof includes an FIR filter according to the invention.

[34] Figure 3 is a functional block diagram of the echo canceller in Figure 2, showing an FIR filter according to the invention together with least mean square elements by which the coefficient for each stage is generated, and including an adjustable delay element.

[35] Figure 4 is a functional block diagram of the 64-delay pipe shown in Figure 3.

[36] Figures 5a and 5b are functional block diagrams showing the FIR filter of Figure 3.

Customer Number: 23624

[37] Figure 6 is a functional block diagram showing the quantizer and downsampling blocks of the FIR filter of Figure 3.

[38] Figure 7 is a flowchart depicting a method of determining how much delay should be provided to the input signal in accordance with the present invention.

[39] Figure 8 is a functional block diagram showing a conventional FIR filter.

[40] Figure 9 is a functional block diagram showing a FIR filter in accordance with a second embodiment of the present invention.

[41] Figure 10 is a functional block diagram showing a FIR filter in accordance with a third embodiment of the present invention.

[42] Figure 11 is a functional block diagram showing an alternate configuration of an FIR filter in accordance with the third embodiment of the present invention.

[43] DETAILED DESCRIPTION OF THE PRESENTLY PREFERRED EMBODIMENTS

[44] First Embodiment

[45] The present invention will now be described with reference with to an echo canceller used in an Ethernet transceiver device. Preferably, the echo canceller is embodied in an Integrated Circuit disposed between a digital interface and an RJ45 analog jack. The Integrated Circuit may be installed inside a PC on the network interface card or the motherboard, or may be installed inside a network switch or router. However, other embodiments include applications in read circuits for disk drives, ghost cancellation in broadcast and cable TV transmission, channel equalization for communication in magnetic recording, echo cancellation, estimation/prediction for speech processing, adaptive noise cancellation, etc. All such embodiments are included within the scope of the appended claims.

[46] While the present invention is described with respect to a digital FIR filter, it is to be understood that the structure and functions described herein are equally

Customer Number: 23624

applicable to an analog FIR. Moreover, while the invention will be described with respect to the functional elements of the FIR filter, the person of ordinary skill in the art will be able to embody such functions in discrete digital or analog circuitry, or as software executed by a general purpose process (CPU) or digital signal processor.

[47] A functional block diagram of an Ethernet transceiver incorporating an FIR filter according to the present invention is depicted in Figure 2. Although only one channel is depicted therein, four parallel channels are typically used in Gigabit Ethernet applications. Only one channel is depicted and described herein for clarity.

[48] A 125 MHz, 250Mbps digital input signal from a PC is PCS-encoded in a PCS encoder 2 and is then supplied to a D/A converter 4 for transmission to the Ethernet cable 6. The PCS-encoded signal is also supplied to a NEXT (Near End Transmitter) noise canceller 8 and to adaptive echo canceller 10. The operation of the echo canceller 10 will be described later herein with respect to Figure 3.

[49] Signals from the Ethernet cable 6 are received at adder 14 and added with correction signals supplied from baseline wander correction block 12 (which corrects for DC offset). The added signals are then converted to digital signals in the A/D converter 16, as controlled by timing and phase-lock-loop block 18. The digital signals from A/D converter 16 are supplied to delay adjustment block 20, which synchronizes the signals in accordance with the four parallel Ethernet channels. The delay-adjusted digital signals are then added with the echo-canceled signals and the NEXT-canceled signals in adder 22.

[50] The added signals are supplied to a Feed Forward Equalizer filter 24 which filters the signal prior to Viterbi trellis decoding in decoder 26. After Viterbi decoding, the output signal is supplied to PCS decoder 28, after which the PCS-decoded signal is supplied to the PC.

[51] The decoder 26 also supplies output signals to a plurality of adaptation blocks schematically depicted at 30 in Figure 2. As is known, such adaptation blocks carry out corrections for such conditions as temperature offset, connector mismatch, etc. The adaptation block 30 provides output to the baseline wander correction circuit 12, the timing and phase-lock-loop circuit 18, the echo canceller

Customer Number: 23624

10, and the NEXT canceller 8.

[52] Each functional block depicted in Figure 2 includes a slave state controller (not shown) for controlling the operation and timing of the corresponding block. A PCS controller 32 controls the slave state controllers of all elements depicted in Figure 2, in a manner to be described below.

[53] Figure 3 is a functional block diagram of the echo canceller 10 shown in Figure 2. In Figure 3, the PCS-encoded logic signal is provided to logic encoder 302 as a five level logic signal (e.g. -1, -0.5, 0, +0.5, +1). The encoder 302 encodes the signal as 3 control bits, which correspond to the five logic levels of the PCS-encoded signal (e.g. -1=100; -0.5=101; 0=010; 0.5=001; 1=000). These control bits are supplied to a first plurality or block of filter stages 304 (comprising taps 0 to 31 of the FIR filter), a second plurality or block of filter stages 306 (comprising taps 32 to 63), a third plurality or block of filter stages 308 (comprising taps 64 to 95), and a fourth plurality or block of filter stages 310 (comprising taps 96 to 127).

[54] Filter blocks 304, 306, 308, and 310 typically have fixed delay periods between each of the taps for constant sampling of the early regions of the input signal where significant signal strength is present. Referring to Figure 1, large amplitudes are present in regions 1 and 2 of the input signal, and (according to the present embodiment) the blocks 304, 306, 308, and 310 receive these regions of the input signal to insure filtering of these significant portions of the signal. A more detailed description of the filter blocks will be provided later herein.

[55] The logic-level-encoded signal from encoder 302 is also supplied to a 64-delay pipe (with 4 increment) 312. The delay pipe 312 is controlled by the echo controller's sequence control state machine 314 so that the portion of the input signal having the most significant echo noise is supplied to filter block 316 for noise cancellation. That is, the region 3 of the input signal is delayed appropriately in delay pipe 64 so that region number 3 is not subjected to echo cancellation (it is "skipped over") until portion 4 can be received and input into filter block 316. This way, not the entire input signal is FIR-filtered, and not as many taps are needed to effectively cancel the echo in the input signal. The method by which

Customer Number: 23624

the signal is selectively delayed will be described in more detail below.

[56] The output of the logic level encoder 302 is also supplied to a quantizer 318 which encodes the three control bits into two logic bits for application to downsampling blocks 322 and 324 (to be described below). For example, the quantizer 318 encodes 000 as 00; 001 as 00; 010 as 10; 101 as 01; and 100 as 01. The quantizer 318 thus performs a rounding function so that the encoded signal may be used to control the least mean squares (LMS) engines 0 through 6.

[57] The LMS engines 4, 5, and 6 are designed to supply tap weighting coefficients to a single block of 32 FIR filter taps, and thus downsampling block 324 can use the same quantizer data for 32 cycles. In contrast, and in accordance with the present invention, LMS engines 0, 1, 2, and 3 are designed to supply tap weighting coefficients to taps 0 to 31 of filter block 304, and downsampling block 322 controls each of these LMS engines in a time-cyclic fashion. This architecture allows more precise filtering of the early regions of the input signal having significant signal strength. For example, at time t_1 , LMS engine 0 supplies a weighting coefficient to tap 0, LMS engine 1 supplies a weighting coefficient to tap 1, LMS engine 2 supplies a weighting coefficient to tap 2, and LMS engine 3 supplies a weighting coefficient to tap 3. At time t_2 , LMS engine 0 supplies a weighting coefficient to tap 1, LMS engine 1 supplies a weighting coefficient to tap 2, LMS engine 2 supplies a weighting coefficient to tap 6, and LMS engine 3 supplies a weighting coefficient to tap 4. In this cyclic fashion, LMS engines 0-3 supply weighting coefficients to more precisely filter the region 1 of the input signal, in contrast to the less precise filtering of the region 2 of the input signal filtered by filter blocks 306, 308, and 310. The above is described in more detail in commonly assigned U.S. Patent application Serial No. 09/465228, filed December 19, 1999 and entitled, "A Method and Apparatus for Digital Near-End Echo / Near-End Crosstalk Cancellation with Adaptive Correlation", the contents of which is incorporated herein by reference.

[58] The quantizer 320 quantizes the output of the delay pipe 312 and supplies it to the downsampling block 324 in a manner similar to that described above with respect to quantizer 318. Downsampling block 326 then controls LMS engine 7 which supplies weighting coefficients to the taps

Customer Number: 23624

128 to 159 of the filter block 316 (which thus filters the adaptively delayed portion of the input signal).

[59] The manner by which the LMS engines generate the tap coefficients will now be described. The LMS engines 0 to 7 input error signals from the FFE 24 or the Viterbi decoder 26 of Figure 2. A memory 330 stores weighting coefficients for each of taps 32-127. As the error signal is received from the FFE 24 or the Viterbi decoder 26, the appropriate coefficients are extracted from memory 330, applied through the corresponding LMS engine, and provided to the appropriate taps 32-127 in order to filter the input signal to eliminate the echo noise in region 2 of the input signal.

[60] In a manner similar to that described above, memory 332 stores coefficients for the taps 0-31 of the filter block 304. The appropriate coefficients are extracted from memory 332 and applied to the appropriate LMS engines 0-3 together with the error signal, and the appropriate coefficients are then supplied to the taps 0-31 to appropriately filter the echo noise in region 1 of the input signal. Similarly, the memory 334 stores coefficients for the taps 128-159, which are selectively applied to the LMS engine 7 together with the error signal. The appropriate tap coefficients are then applied to filter block 316.

[61] Figure 4 is a functional block diagram of the 64-delay element 312 of Figure 3. As can be seen, the 64 delay elements are grouped in sets of four delay elements 412, 414, 416, and 418. The logic level-encoded signal S is input to the delay pipe and may be delayed in increments of four by activation of control signals at gates 420, 422, and 424. The control signals are supplied by the sequence control state machine 314, and are varied in accordance with which portion of the input signal is to be skipped, as will be described below.

[62] Figure 5a is a functional block diagram of the FIR filter showing how the variable delay D is supplied to an existing delay element 512 in order to variably adjust the input signal to skip the desired portion thereof. In Figure 5a, the logic level-encoded signal S is supplied, for example, to a first element 520 having a time delay t_1 . A tap coefficient C_0 is applied to a multiplier 505 in order to weight the first tap of the FIR filter. The weighted signal is then provided to a summer 515 where it is added to the outputs of the other stages (to be

Customer Number: 23624

described below), and then output as signal S_0 . The signal S is also supplied to the multiplier 518 for multiplication by coefficient C_1 , and addition with the other outputs at summer 514. Of course, any number of additional stages like 520 may be provided prior to the output, as required.

[63] The input signal S is also supplied to delay element 512 having a variable delay D . The thus-delayed signal S_{vd} is then provided to a series of sequential delay elements including delay element 506, which preferably also has a fixed delay time t_1 . The delayed signal S_{vd} is also supplied to multiplier 516 for multiplication by coefficient C_{n-2} and addition in summer 513, as shown. The output of delay element 506 $S_{vd}+t_1$ is supplied to both another delay element 502 (having a t_1 delay) and a multiplier 510 where it is multiplied by coefficient C_{n-1} .

The output of element 502 $S_{vd}+t_1+t_1$ is supplied to multiplier 504 where it is multiplied by coefficient C_n and then added, in adder 508, to the output from multiplier 510. In this manner, the series of weighted tap coefficients and corresponding input signals are processed through the FIR filter, in a manner known to those of skill in the art.

[64] The appropriate number of stages with corresponding delay elements are provided in order to properly filter the regions of the input signal having significant signal strength, such as regions 1 and 2 in Figure 1. However, to skip those insignificant portions of the signal (such as region 3), the element 512 is provided with the variable delay D in accordance with control signal C_t supplied from the sequence control state machine 314. According to the present invention, the variable delay D may be selected to skip any portion of the input signal which is not to be filtered. Preferably, a later portion of the input signal will be filtered since significant echo typically resides therein. Accordingly, after element 512, any number of additional stages like elements 502 and 506 are provided, typically having the same fixed time delay t_1 . The number of additional stages after stage 512 may be varied to appropriately filter the echo regions of the input signal.

[65] Figure 5b shows an alternative wherein the delay element 584 is provided to the undelayed portion of the input signal S to skip portions thereof. Like reference numerals represent like structure. In Figure 5b, the input signal S is supplied to both of multipliers 590 and 592

Customer Number: 23624

where it is respectively multiplied by coefficients C_0 and C_1 . The delayed signal S_{vd} output from element 584 is, after any number of intervening stages, supplied to both multipliers 510 and 504 where it is respectively multiplied by coefficients C_{n-1} and C_n . The output of multiplier 504 is delayed in a delay element 502 having a t_1 delay, and then supplied to adder 508 where it is added to the output from multiplier 510. The output of adder 508 is then supplied to a delay element 506 having a delay of t_1 , and the output of 506 is, in turn, provided (after any number of intermediate stages) to the adder 514 where it is added with the output of multiplier 590. The output of adder 514 is provided to a delay element 586 having a t_1 delay. The output of the element 586 is added, in adder 588, to the output of multiplier 592, and the output of adder 588 is the output signal S_0 .

[66] In a further alternative to the above arrangement, variable delays may be provided to more than one filter block. For example, filter block(s) 310 and/or 308 and/or 306 may also be supplied with variable delays so that any portions of the input signal may be skipped or filtered as the circuit designer requires. All such alternatives are included within the scope of the appended claims.

[67] Figure 6 is a functional block diagram of the quantizer and downsampling circuits of Figure 3. The quantizer 318 receives the logical level-encoded signal S from the input of delay pipe 312. The output of quantizer 318 is provided to both the downsampling block 324 and a multiplexer 612. The multiplexer 612 outputs the quantizer signal to a one-cycle delay element 614, which supplies the down-sampled signal to LMS engine 3. In a similar manner, delay elements 616, 618, and 620 respectively provide down-sampled signals to LMS engines 2, 1, and 0, after the appropriate delay. The output of delay element 620 is also returned to the multiplexer 612, as shown.

[68] The output of downsampling block 324 is provided to the LMS engines 6, 5, and 4, as was described above with reference to Figure 3. Also, the output of the delay pipe 312 is supplied to the quantizer 320 which supplies the downsampling block 326 and LMS engine 7, as shown.

[69] In operation, those portions of the input signal which may be skipped by the FIR filter must first be determined. Preferably, this is done by injecting a test signal into the Ethernet cable and then receiving the

Customer Number: 23624

return signal, such as the waveform depicted in Figure 1. However, the procedure for determining the insignificant portions of the input signal may be performed at any convenient time, such as when the Ethernet is first powered on, after any Ethernet device has been plugged into the network or unplugged from the network, during any lull in Ethernet communications, on a periodic basis, or continually. The signal used to determine the delay may also be any appropriate signal such as a test signal, a series of test signals, or by using actual Ethernet communication signals on-the-fly.

[70] The method of determining how much delay to be supplied to the input signal in accordance with the embodiment of Figure 3 will now be described with respect to the flow chart of Figure 7. This process is preferably carried out within the sequence control state machine 314, although any convenient processor and memory may be used. In Figure 7, when the Ethernet is first powered-up, data starts to be supplied to the Ethernet cable 6 at step S1. At step S2, the return signal is received and then filtered in the FIR filter using blocks 304, 306, 308, 310, and 316 contiguously so as to filter a continuous portion of the return signal. At step S4, it is determined which tap of taps 128-159 has received the maximum return signal strength. This tap is labeled tapmaxd. At step S5, tapmaxd is compared with the stored tapmaxs, and the tap having the maximum signal strength is then stored as the new tapmaxs. Of course, for the first determination, the initial tapmaxd will be stored as tapmaxs. In order to avoid storing unexpectedly large signal strength caused by noise, multiple looping for comparison is preferably employed. For example, if 32 taps are compared and tap 7 is identified as tapmaxs, the comparison will be repeated multiple times. Every comparison, tap 7 will be replaced with tapmacxs even though the tapmaxs is larger than tap 7, in order to avoid a lock up error.

[71] At step S6, it is determined whether the end of the return signal has been reached. If the end of the return signal has not been reached, the process proceeds to step S7 where a 32 tap delay is applied to skip a portion of the return signal. Of course, any amount of tap delay (1 tap, 4 taps, 8 taps, 16 taps, 64 taps, etc.) may be used in any combination by the circuit designer to flexibly configure the FIR filter. The process then returns to step S4 to determine which tap of the newly-filtered signals has the maximum signal strength. Again, the determined tapmaxd is compared with the stored tapmaxs, and the

Customer Number: 23624

maximum value is stored as the new tapmaxs in step S5.

[72] One algorithm for performing steps S4, S5, S6, and S8 of Fig. 7 is as follows:

[73] Let n = the number of stages in the FIR filter.
[74] Let tap[i] = the ith stage of the FIR filter.
[75] Let {tap[i]} = the coefficient value of the ith stage of the FIR filter.
[76] Let Maxcoeff = the absolute value of the maximum coefficient value in the FIR filter.
[77] Let m = the index of which tap coefficient is written into Maxcoeff.

[78] At time = 0,
[79] Maxcoeff • {tap[0]}
[80] m • 0

[81] At time = i, (where i > 0, i.e., 1, 2, 3, 4,...)
[82] if (en_search) //where en_search enables the search for Maxcoeff
[83] begin
[84] if (Maxcoeff <•{tap[i]}• or m =i)
[85] begin
[86] Maxcoeff • •{tap[i]}•
[87] m • I
[88] end
[89] else
[90] begin
[91] Maxcoeff • Maxcoeff
[92] m • m
[93] end
[94] end
[95] else
[96] begin
[97] Maxcoeff • Maxcoeff
[98] m • m
[99] end.

[100] In this iterative manner, the last filter block 316 is successively moved across the later portions of the return signal identifying which portion(s) of the return signal have the maximum signal strength. When the filter block 316 has reached the end of the return signal, step S8 is performed wherein the stored tapmaxs is set as the

Customer Number: 23624

center tap of the filter block 316. Now, the filter block 316 will be applied to the center of the later portion of the return signal having the most significant signal strength. The required delay may be determined algorithmically or from accessing an entry from a lookup table. The delay required to so-position filter block 316 is then stored in the memory of sequence control state machine 314 so that all Ethernet signals received from the Ethernet cable 6 may be FIR-filtered in accordance with the thus-configured filter blocks to skip those portions of the signal having insignificant signal strength, while filtering the remaining signal. In such a manner, Ethernet signals typically requiring more than 220 taps for proper FIR filtration can be adequately filtered with an FIR filter having only 160 taps.

[101] Thus, what has been described is method and apparatus for controlling an FIR filter so as to delay the input signal to skip over portions of that signal having insignificant signal strength. This allows the FIR filter to have fewer taps, consuming less power and less space on the Integrated Circuit.

[102] The individual components shown in outline or designated by blocks in the attached Drawings are all well-known in the FIR filtering arts, and their specific construction and operation are not critical to the operation or best mode for carrying out the invention.

[103] SECOND EMBODIMENT

[104] Figure 8 is a block diagram of a conventional FIR filter. As shown therein, input data is applied to one input of multiplier 82-1 to be multiplied by a first coefficient supplied from coefficient generator or preferably LMS engine 50. The input is applied to delay circuit 84-2 of the next stage and the output of multiplier 82-1 is supplied adder 86-2. The output of delay circuit 84-2 is applied to one input of multiplier 82-2 and to delay circuit 84-3 of the next stage. Multiplier 82-2 multiplies the output of delay circuit 84-2 by a second coefficient, which is also supplied by LMS engine 50. The coefficients supplied to their respective multipliers each contain a plurality of bits. In the preferred embodiment of the present invention each coefficient is 13 bits. LMS engine 52 supplies the coefficients to the multipliers by respective wirings. In the preferred embodiment each coefficient requires 13 conductors per wiring or a total of 2080 conductors for

Customer Number: 23624

160 taps. LMS engine 50 also supplies the coefficients to memory 52 at a higher resolution, which in the preferred embodiment is 20 bits. An output of memory 52 is fed back to LMS engine 50 for further calculations. The output of multiplier 82-1 is added to the output of multiplier 82-2 by adder 86-2. The succeeding stages are similarly configured. An FIR filter having 2080 conductors is more complex and consumes a significant amount of area which results in a larger die size.

[105] Figure 9 is a block diagram of the FIR filter in accordance with the second embodiment of the present invention. The second embodiment overcomes the above-discussed problem by sharing the wirings for all the coefficients supplied from LMS engine 50 to its corresponding tap of the FIR filter. Wirings are formed from a conductive material, such as by way of example aluminum, copper, polysilicon and the like. Referring to Figure 9, LMS 50 supplies each of the coefficients via a shared or common set of wirings to a respective memory (80-1...80n) for each corresponding tap. LMS engine 50 and memories (80-1...80-n) are under the control of controller 55. Memories 80-1..80n are preferably implemented as latches. As would be appreciated by one of ordinary skill in the art, other appropriate circuitry may be utilized, such as flip-flops, SRAM, DRAM, and the like. Controller 55 sequentially selects the coefficient to be provided by LMS engine 50 and a respective memory (80-1...80-n) to store the coefficient. The stored coefficient is then provided to a corresponding multiplier (82-1..82-n) to perform the multiplication operation. As used herein for this embodiment, the term LMS engine shall include individual LMS circuits to generate coefficients for each tap or an LMS circuit to generate coefficients for a group of taps, or any combination thereof. In the preferred embodiment the coefficient wiring requires 13 conductors and the number of taps is 160. Therefore the second embodiment of the present invention requires 13 conductors for the shared coefficient wiring and one control conductor for each tap or 160 conductors. In other words in the preferred embodiment 173 conductors are required. The delay times of delay circuits 84-2 ...84-n may be equal or some of the delay circuits may set to different values in accordance with the first embodiment of the present invention.

[106] THIRD EMBODIMENT

[107] Reference is now made to Figure 10, which illustrates a block diagram of the third embodiment in accordance with the present invention. As shown therein, the third

Customer Number: 23624

embodiment is similar to the second embodiment except the third embodiment comprises a selector circuit (which is comprised by a combination shift register 120 and multiplexer 122) to locally generate the control signals for controlling the memories 80-1...80-n in synchronization with the coefficients output by LMS engine 50. The number of registers in shift register 120 equals the number of taps. In the preferred embodiment there are 160 taps and 160 registers in shift register 120. The operation of the third embodiment is as follows. Controller 50 generates an initialization signal for LMS engine 50 and multiplexer 122. At that time LMS engine 50 outputs a first coefficient and at each subsequent clock signal outputs a successive coefficient. Upon receiving the initialization signal, multiplexer 122 selects the first input (value = 1) and loads the "1" into the first register of shift register 120. The first register corresponds to memory 82-1 of the first tap, and the first coefficient is stored therein. In response to the clock signal, the "1" is shifted in shift register 122, so that subsequent memories are enabled in synchronization when its corresponding coefficient is output by LMS engine 50. In the third embodiment, the number of conductors is equal to the width of the shared coefficient and one conductor for the initialization signal from controller 55. In the third embodiment the number of conductors is 13 + 1 or 14.

[108] Figure 11 shows an arrangement in which one LMS engine is provided for each 32 taps of the FIR filter. More specifically, the FIR filter comprises five FIR filter sections 200-1...200-5, each having 32 taps. The coefficients of FIR filter sections 200-1...200-5 are supplied from LMS engines 50-1...50-5. As can be seen from Figure 11, each FIR filter sections requires 14 conductors (13 conductors from LMS engine 50-n and one from controller 55-n). Thus an FIR filter having 160 taps arranged in five FIR filter sections requires 70 conductors.

[109] While the present invention has been described with respect to what is presently considered to be the preferred embodiments, it is to be understood that the invention is not limited to the disclosed embodiments. To the contrary, the invention is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims. The scope of the following claims is to be accorded the broadest interpretation so as to encompass all such modifications and equivalent structures and functions.